AD-A103 005 YALE UNIV NEW HAVEN CT DEPT OF COMPUTER SCIENCE F/G 9/2
WHAT'S THE POINT?(U)
MAY 81 R C SCHANK, G C COLLINS, E DAVIS N00014-75-C-1111
UNCLASSIFIED TR-205 NL

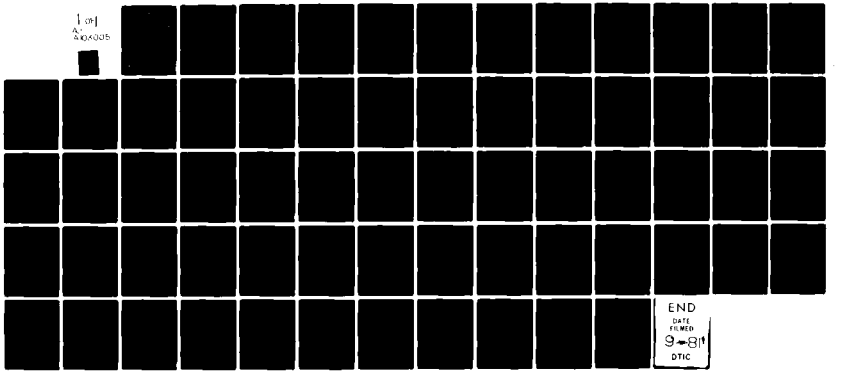

What's the Point?

by

Roger C. Schank, Gregg C. Collins, Ernest Davis
Peter N. Johnson, Steve Lytinen, and Brian J. Reiser

Research Report #205
May 1981

YALE UNIVERSITY
DEPARTMENT OF COMPUTER SCIENCE

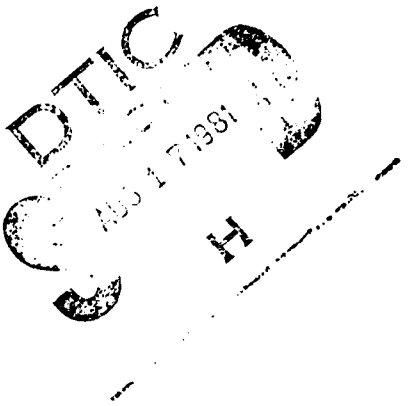

What's the Point?

by

Roger C. Schank, Gregg C. Collins, Ernest Davis,

Peter N. Johnson, Steve Lytinen, and Brian J. Reiser

Abstract

Understanding dialogue usually requires determining the intent, or point, of the utterance. Finding the point serves to constrain further processing. We present a categorization of points, and we propose algorithms and heuristics for deriving the point of a given utterance.

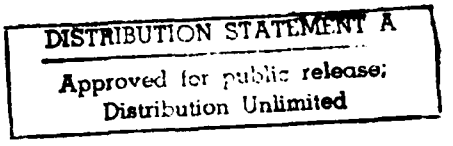

The research described here was done at the Yale Artificial Intelligence Project and is funded in part by the Office of Naval Research under contract N00014-75-C-1111 and the National Science Foundation under contract IST 7918463.

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

REPORT DOCUMENTATION PAGE

READ INSTRUCTIONS BEFORE COMPLETING FORM

1. REPORT NUMBER # 205	2. GOVT ACCESSION NO. AD-A103005
3. RECIPIENT'S CATALOG NUMBER	
4. TITLE (and Subtitle) What's the Point?	5. TYPE OF REPORT & PERIOD COVERED Technical Report
	6. PERFORMING ORG. REPORT NUMBER
7. AUTHOR(s) Roger C. Schank, Gregg C. Collins, Ernest Davis Peter N. Johnson, Steve Lytinen, and Brian J. Reiser	8. CONTRACT OR GRANT NUMBER(s) N00014-75-C-1111 IST7918463
9. PERFORMING ORGANIZATION NAME AND ADDRESS Yale University - Computer Science Department 10 Hillhouse Avenue New Haven, Connecticut 06520	10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS
11. CONTROLLING OFFICE NAME AND ADDRESS Advanced Research Projects Agency 1400 Wilson Boulevard Arlington, Virginia 22209	12. REPORT DATE May 1981
	13. NUMBER OF PAGES 69
14. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office) Office of Naval Research Information Systems Program Arlington, Virginia 22217	15. SECURITY CLASS. (of this report) unclassified
	15a. DECLASSIFICATION/DOWNGRADING SCHEDULE

16. DISTRIBUTION STATEMENT (of this Report)

DISTRIBUTION STATEMENT A
Approved for public release;
Distribution Unlimited

DTIC
JUL 17 1981

17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)

18. SUPPLEMENTARY NOTES

19. KEY WORDS (Continue on reverse side if necessary and identify by block number)

Artifical Intelligence
Natural Language Processing
Dialog
Goal-Based-Processing
Knowledge Representat

20. ABSTRACT (Continue on reverse side if necessary and identify by block number)

Understanding dialogue usually requires determining the intent, or point, of the utterance. Finding the point serves to constrain further processing. We present a categorization of points, and we propose algorithms and heuristics for deriving the point of a given utterance.

DD FORM 1 JAN 73 1473 EDITION OF 1 NOV 65 IS OBSOLETE
S/N 0102-LF-014-6601

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

-- OFFICIAL DISTIRUBTION LIST --

Defense Documentation Center — 12 copies
Cameron Station
Alexandria, Virginia 22314

Office of Naval Research — 2 copies
Information Systems Program
Code 437
Arlington, Virginia 22217

Dr. Judith Daly — 3 copies
Advanced Research Projects Agency
Cybernetics Technology Office
1400 Wilson Boulevard
Arlington, Virginia 22209

Office of Naval Research — 1 copy
Branch Office - Boston
495 Summer Street
Boston, Massachusetts 02210

Office of Naval Research — 1 copy
Branch Office - Chicago
536 South Clark Street
Chicago, Illinois 60615

Office of Naval Research — 1 copy
Branch Office - Pasadena
1030 East Green Street
Pasadena, California 91106

Mr. Steven Wong — 1 copy
New York Area Office
715 Broadway - 5th Floor
New York, New York 10003

Naval Research Laboratory — 6 copies
Technical Information Division
Code 2627
Washington, D.C. 20375

Dr. A.L. Slafkosky — 1 copy
Commandant of the Marine Corps
Code RD-1
Washington, D.C. 20380

Office of Naval Research — 1 copy
Code 455
Arlington, Virginia 22217

Office of Naval Research — 1 copy
Code 458
Arlington, Virginia 22217

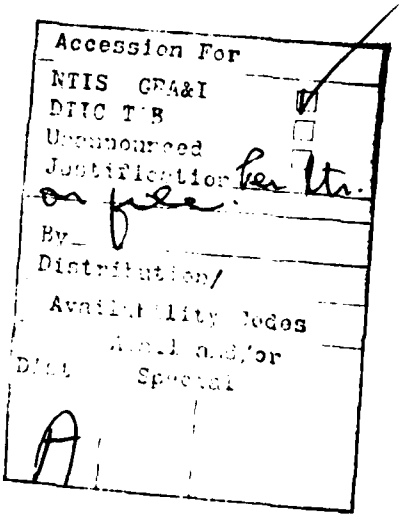

Naval Electronics Laboratory Center Advanced Software Technology Division Code 5200 San Diego, California 92152	1 copy
Mr. E.H. Gleissner Naval Ship Research and Development Computation and Mathematics Department Bethesda, Maryland 20084	1 copy
Captain Grace M. Hopper, USNR Naval Data Automation Command, Code 00H Washington Navy Yard Washington, D.C. 20374	1 copy
Dr. Robert Engelmore Advanced Research Project Agency Information Processing Techniques 1400 Wilson Boulevard Arlington, Virginia 22209	2 copies
Professor Omar Wing Columbia University in the City of New York Department of Electrical Engineering and Computer Science New York, New York 10027	1 copy
Office of Naval Research Assistant Chief for Technology Code 200 Arlington, Virginia 22217	1 copy
Captain Richard L. Martin, USN Commanding Officer USS Francis Marion (LPA-249) FPO New York 09501	1 copy
Major J.P. Pennell Headquarters, Marine Corp. (Attn: Code CCA-40) Washington, D.C. 20380	1 copy
Computer Systems Management, Inc. 1300 Wilson Boulevard, Suite 102 Arlington, Virginia 22209	5 copies
Ms. Robin Dillard Naval Ocean Systems Center C2 Information Processing Branch (Code 8242) 271 Catalina Boulevard San Diego, California 92152	1 copy

Dr. William Woods — 1 copy
BBN
50 Moulton Street
Cambridge, MA 02138

Professor Van Dam — 1 copy
Dept. of Computer Science
Brown University
Providence, RI 02912

Professor Eugene Charniak — 1 copy
Dept. of Computer Science
Brown University
Providence, RI 02912

Professor Robert Wilensky — 1 copy
Univ. of California
Elec. Engr. and Computer Science
Berkeley, CA 94707

Professor Allen Newell — 1 copy
Dept. of Computer Science
Carnegie-Mellon University
Schenley Park
Pittsburgh, PA 15213

Professor David Waltz — 1 copy
Univ. of Ill at Urbana-Champaign
Coordinated Science Lab
Urbana, IL 61801

Professor Patrick Winston — 1 copy
MIT
545 Technology Square
Cambridge, MA 02139

Professor Marvin Minsky — 1 copy
MIT
545 Technology Square
Cambridge, MA 02139

Professor Negroponte — 1 copy
MIT
545 Technology Square
Cambridge, MA 02139

Professor Jerome Feldman — 1 copy
Univ. of Rochester
Dept. of Computer Science
Rochester, NY 14627

Professor Nils Nilsson — 1 copy
Stanford Research Institute
Menlo Park, CA 94025

Dr. Alan Meyrowitz — 1 copy
Office of Naval Research
Code 437
800 N. Quincy Street
Arlington, VA 22217

Dr. Edward Shortliffe — 1 copy
Stanford University
MYCIN Project TC-117
Stanford Univ. Medical Center
Stanford, CA 94305

Dr. Douglas Lenat — 1 copy
Stanford University
Computer Science Department
Stanford, CA 94305

Dr. M.C. Harrison — 1 copy
Courant Institute Mathematical Science
New York University
New York, NY 10012

Dr. Morgan — 1 copy
University of Pennsylvania
Dept. of Computer Science & Info. Sci.
Philadelphia, PA 19104

Table of Contents

1. Introduction

Part of understanding involves assessing why a speaker says what he says. Understanding what someone has said involves understanding what the speaker intended the hearer to do with what he said.

It is quite possible for someone to understand exactly what a speaker said, but not understand why it was said, or why it was said to him in particular. Alternatively, an understander may think he knows why something was said to him, but he may be incorrect. In such cases, we usually say he has "missed the point", and the speaker may find a response he receives to be incoherent.

A very important part of what we do when we understand is attempt to figure out the point of what we are being told. This search for the point guides processing, and thus directs the inferencing process. We make choices early on in processing that cause us to go one way or another in what aspects of an input we dwell on and which we ignore. Clearly this initial assessment of what is important is a key element in the understanding process and in the subsequent memory organization that results from the way we have processed an input.

Recent research in natural language has focused on the examination of the types of mental knowledge structures necessary to build a conceptual representation of a discourse. Scripts, plans and goals have been proposed [Schank and Abelson 77] and both embodied in programs, [Carbonell 79]; [Cullingford 78]; [Wilensky 78] and tested in psychological experiments (see [Abelson 80] for a review). However, as

[Schank 78] pointed out, an understanding system utilizing inferences driven by these types of knowledge structures faces the problem of controlling those inferences. That is, one cannot constantly make all the causal, scriptal, and goal-related inferences possible on every input, or a combinatorial explosion would result. There must be some guidance in the application of these knowledge-based inference rules. [Schank 78] suggested that the understander would focus attention on those concepts that were interesting at the expense of other less interesting concepts. Interestingness was defined in terms of several absolute categories (death, sex, etc.) and some relative rules which are based on deviations from expectations.

One way we can know where to begin in processing is by attempting to assess the point of what we have heard. To make this clear, consider some examples of what we mean by points.

Suppose you are told that the Zambians are going to invade your country tomorrow. What would the point of your being told that be? Clearly there is no one correct answer to such a question. A person could have any number of points in mind that a particular utterance expresses. For this sentence, some possible points are:

1. Your're going to be in trouble very soon.
2. Boy, this is really the last straw. Things couldn't get any worse for me.
3. Get your guns right now so we can get ready to repel the invaders.
4. That explains why the their ambassador decided to leave yesterday.

5. Head for the hills immediately if you want to save your life.

6. My assertion that invasions only happen on Tuesdays is again shown to be correct.

7. Who would have thought that they would really invade?

It may seem here that almost anything can be a point. In one sense this is true. Nearly any proposition can realize a point that someone is trying to express. But, the notion of a point is one that does not depend upon the content of a proposition directly. Rather whether something is a point depends upon its relationship to what surrounds it. Saying that a given proposition is a point is a statement about its role with respect to what you as an understander are supposed to do with that proposition. Viewed from a processing perspective, when a speaker makes a point, he is trying to direct the understanding process of the listener. In this sense, points provide a much needed focussing mechanism during understanding. Thus tracking the point is another way to limit the number of inferences in the understanding process.

The key question is then, in what different directions can one proceed while processing a proposition? After that is known, then the question that follows is: How does one recognize which of these understanding directions the speaker intended in making his point? And finally, once the speaker's point has been determined, how does this affect the listener's subsequent understanding process?

We define a point as being a proposition derivable from an input that conforms to one of the following seven statements:

1. This new information that I just gave you will affect you.

2. Here is a problem of mine, you should understand my feelings.

3. Here is a problem of mine that I need help on.

4. This new information should help explain things.

5. This is what I am saying you must do.

6. Therefore, X is correct.

7. Here is something that is true that should interest you.

The point here, if we may use the word, is that when someone has a point, that is, when some statement is intended to express a point, one of the above seven points is what is being expressed. When we are confused by something someone says and ask 'whats the point?' we mean to convey the fact that we are having trouble deriving one of the seven points above from what was said.

These seven points relate to the understanding process in the following way. As each statement is processed, we attempt to derive the speaker's point as one of these seven. Then, knowing what the speaker intended, we may elect to process his statement to a deeper level in accordance with his point. This is accomplished by specific procedures associated with the point categories. However, as we shall see later, recognizing the speaker's intended processing category and actually processing the statement in keeping with his intentions are two different things. In other words, a listener may recognize the speaker's point, but still not process the statement in the targeted category. For now, consider this first task of deriving the speaker's point.

To derive a point from what we hear or read, we must constantly be asking ourselves these questions that relate to the points above:

1. How does this affect me?
2. Do I empathize with this problem?
3. How can I help here?
4. Is this new and relevant information?
5. Is the speaker telling me that I must do this.
6. How does what I just heard relate to a prior point of view that this person expressed?
7. Does this interest me?

In asking ourselves these questions, we determine how the speaker intended us to process his statement. From there, we may choose to process the statement deeper in keeping with these processing categories.

Thus the problem is:

- What kinds of points are there?
- How are points derived?
- How does subsequent processing proceed after the speaker's point has been determined?

We shall now proceed to discuss these issues.

1.1 Kinds of Points

Affective Points

People usually have a set of plans and goals that are presently active in their minds. Some of these goals may be long-term, such as the goal of becoming President, or some may be very short-term, such as the goal of finding a pencil. Often, the point of a speaker's statement

is to address one or more of the understander's plans and goals, or to call to the understander's attention that some piece of information has some positive or negative affect on one of the understander's plans or goals. These points are called affective points.

An example of an affective point was given in the Zambian example in the previous section. In the case in which the speaker meant, "Head for the hills immediately if you want to save your life," he was obviously addressing one of the understander's goals. After understanding that this was what was meant, the understander would not ask, "What's the point?"

A point can also be affective if it addresses the plans and goals of someone which the understander has some emotional feelings toward. Thus, if your sister lived in Peoria and someone told you, "The Zambians are invading Peoria!", his point would be affective. Again, you would not ask, "what's the point?" These emotional feelings need not be positive, however. If your worst enemy lived in Peoria and someone told you "The Zambians are invading Peoria!", you still could very well derive an affective point.

Empathetic Points

Conversations between good friends often deal indirectly with the feelings of the conversants. What is the point in such conversations? The speaker may not be conveying anything which will directly affect the listener's life, although the expression of an emotional or problem situation is certainly of interest to a friend. Indeed, part of the meaning of a friend is that things that concern one's friends are of

interest to oneself. Yet the expression of merely interesting items does not seem to capture the importance of such dialogs.

If one friend worriedly says to another, "Oh, have my wife and I been fighting a lot lately!", he would probably not respond, "So, what's the point?" Clearly, the speaker has in mind a particular concept that he wants to communicate, and a very particular way in which he wants his friend to process that concept. At the lowest level, he wants the listener to understand that he is having some sort of marital problems. In this case, however, the speaker not only wants his friend to understand the specific content of the experiences he is describing, but to understand how he is feeling about them. There may be no request for help, or implication of any personal effect on the listener. Instead, what the speaker desires is for the listener to process the input almost as if he might be the focus of the situation, so as to understand the speaker's point of view. When the speaker's point in communicating is getting someone else to understand how he is feeling about a particular situation, we call this an <u>empathetic point</u>. He is saying, in essence, "Here is something that is affecting me, and I just want you to understand how I feel."

<u>Need Points</u>

Someone else's problem may not be your problem but it most certainly can be his point. This was illustrated by the example above in which the worried friend describes his frustration in having an unhappy marriage. However, in many cases when a person describes his problems, he is seeking much more than an empathetic understanding from

the listener. He is often expressing his needs with the hope of getting help.

When someone tells you that his car has run out of gasoline, he is probably not making an empathetic point. He may well be asking for help. Asking for help is his point. Here again, it is bizarre to say, "what's your point?" When we are asked for help, either directly or indirectly, the point being made is simply that you should attempt to provide that help. As another illustration, consider the case in which one student tells another about his difficulty in understanding algebra. He is probably seeking some tutoring, particularly if the other student is known to do well in mathematics courses.

One question that comes to mind here is, if he wanted tutoring or guidance, then why didn't he simply ask for it? This speaks to the heart of the issue of points. Points are frequently expressed indirectly. One reason for this is that social rules often prohibit direct expression of ideas and desires. If we ask for something indirectly, we are not so devastated when we are refused. In this example, by not asking for help directly, the student has left his prospective tutor with the option to politely avoid the issue of providing assistance.

Explanatory Points

As understanders of the world we seek to know why things happen the way they do. Thus, we are often in the situation of not understanding something and wishing we did. Speakers often help us by making explanatory points, which fill the gaps in our understanding. For

example, in the exchange

A: Why are the banks closed today?

B: Lincoln's birthday.

B is making an explanatory point. He provides a cause for a fact which A finds unexplained and anomalous. Teaching consists, in great measure, of presenting explanatory points.

An important class of explanatory points are those which explain or defend the speaker's behavior (or, by extension, the behavior of someone with whom he identifies himself). For example, if a teacher asks, "Why are you late for class?" a student may answer "I had to go to the dentist." The notable point about such examples is that very likely the teacher does not care why the student is late, as long as there is some acceptable reason. It is a social rule that, if one is obliged to behave in an unusual or offensive manner, one will provide an explanation.

Prescriptive Points

Our world is full of rules. Many of these rules can be broken down into explanatory points. When we tell a child "don't play with matches," this admonition can be backed up with an explanatory point concerning possible cause and effect relationships concerning the possible dangers of fire. However, not all rules are thus decomposable. In particular, religious and moral rules can often be broken down no further - ultimately we believe they should be obeyed either because some higher authority has prescribed it or because they just seem to be right, depending on our particular religion/philosophy.

Furthermore, rules which could in theory be broken down are often

not, particularly with children who may be incapable of understanding the actual causal points involved, or by parents whose desire to explain has reached a limit. In reality, children are in fact often given rules for what to do and what not to do without any real explanation being provided.

Many rules are not broken down because whatever causal notions underlie them are too poorly understood to be specified. We often pass on information like "hitting the TV on the side with your shoe seems to improve the picture". This, too, is prescriptive. This information does not pretend to be an explanation of anything, it is merely advice on how to succeed in some of your goals.

Because of this, a lot of the input a person gets is in the form of prescription (or proscription) alone. In fact, in comparison to explanatory points, which have the same ultimate goal of improving the performance of the individual, the number of prescriptive points is probably far greater. Whole disciplines in which explanations have become extremely complex and convoluted often turn more and more to these types of points, which exhibit themselves as superstitious behavior. In computer science, for example, programming, and in particular debugging, are more often done by ritual than by invoking a clear understanding of the situation.

It is clear that the process of internalizing a prescription is different from any of the other types of processing we have examined. Rather than being linked into memory in a causal manner, as explanations

are, prescriptions are simply connected to whatever goals and plans they are intended to govern, serving as constraints on later instances of the planning process.

Argument Points

When we get into an argument, it is often our fervent desire to show our opponent that he is wrong and we are right. In such situations the point of our remarks may well be just that: that our position is right, or our opponent's is wrong. Any statement which is meant to supply evidence for some proposition has implicitly as its point "therefore (that proposition)". Thus, while attempting to prove that the Arabs are the aggressors in the Middle East, an Israeli's point is that that is true.

Of course, this situation will be complicated somewhat by the fact that arguments often proceed by spawning subpropositions which provide evidence reflecting on the top level proposition being argued. Thus when an embedded point of evidence occurs in an argument, it is somewhat ambiguous whether the point is that the nearest subproposition is correct, or that the top level proposition is correct (or incorrect), or both. Clearly the situation can become complex, given several levels of propositions being debated. An argument has a life of its own, and often functions in a manner quite unrelated to normal conversation. Therefore, the concept of points may not be the most natural way to view the situation in an argument [Flowers, McGuire, and Birnbaum 81].

However, argument points are not necessarily limited to arguments. It is not at all uncommon to make a statement in conversation with the

point of supplying evidence for a previous statement. For example, one might say, "The younger generation is totally degenerate. Judy and Richard are living together!" The second statement has the point of providing evidence for the first. Note that this accounts for the fact that a reaction based on affect or anterestingness of the second statement would "miss the point" slightly, and if the speaker cared to pursue it he would have to redirect the listeners attention to his opinion of the younger generation.

Interestingly, it is quite possible that someone would say "Judy and Richard are living together" with the same meaning as the statement in the above example. In other words, it could be implicit that the speaker means to point out that the younger generation is degenerate. In this instance, we face something of a dilemma, since the statement is still an argument point, in that it provides evidence for another proposition, but it also serves to introduce that proposition which presumably is intended as an affective point. It seems clear that the main point is whatever the point of the implicit proposition is, and that the argument point is secondary.

Interest Points

Much of what one hears is told merely for the interest value. If someone tells you that the Yankees are trying to sign Dave Winfield, he is (in most cases) stating a fact that is valuable for you to know merely out of idiosyncratic interests, rather than out of any actual relevance to yourself or himself. The ideal response, from the speaker's point of view, is an expression of interest or excitement on your part:

"Really?" or "That's wonderful" or "How do you know?". This can go wrong in several ways. You may, in fact, be interested in baseball, but have a very low opinion of Dave Winfield; in which case you may respond "So what? It won't do them any good." You may have no interest in baseball, in which case you will generally tell the speaker so: "I'm sorry, I don't like baseball". Or the speaker may be right as to your interest, but mistake your knowledge state, leading to the responses "Who is Dave Winfield?" or "Who are the Yankees?"

The example with Dave Winfield is based on an interest idiosyncratic to the individual hearer. Other statements, like "Someone shot the President" may be assumed to be interesting to most hearers. Criteria for interestingness are sketched in [Schank 78].

1.2 Comparison to Other Work

Morgan [Morgan 78] arrives at conclusions which are very similar to ours. He states:

> From [the content of the speech acts] H [the hearer] must infer what modifications he is intended to make in his model of the world ... For each speech act H must infer a purpose that is consistent with the purposes he inferred for earlier speech acts; or he must revise earlier hypotheses about purposes accordingly. Questions H must infer answers to are "Why did the speaker perform this particular speech act at this particular point in the text?" and "Why does he want me to have this particular fact just now?" ... From speech acts and their purposes taken jointly, he must construct a hypothesis of the speaker's goal in the test, and of the plan that the speaker is following in advancing toward that goal ...

Our overall model is in close agreement with Morgan's. There are, however, many important aspects of the problem which are not discussed in his brief paper. In particular, we will deal with the categorization

of speakers' points; the use of non-linguistic context to aid in determining the speaker's point; and the interactions between determining the speaker's point, and determining the propositional content of speech acts.

Our theory, and Morgan's, clearly lie within the philosophy of "language as use", of which the chief proponents are Grice [Grice 75], Austin [Austin 62], and Searle [Searle 69]. Grice presents four rules which govern conversational speech, and he gives a number of examples where the point of an utterance may be found by assuming that the speaker follows these rules. Grice suggests that these rules can probably be derived from an analysis of the goals of communication, but gives no indication of how this may be done. Our analysis begins at the opposite end. We consider the general problem of describing and determining a speaker's goals, and we see, in many cases, how a Gricean analysis may be performed.

[Austin 62] and [Searle 69] are likewise related. The object of their research, however, is quite different. Austin and Searle are interested in speech acts as acts which do something like "asserting truths about the world" or "commit the speaker to a future action" or "attempt to get the hearer to do something", etc. Our primary interest is in communication, viewing communication as an attempt by the speaker to have the hearer perform more or less specific mental operations, like "form a causal chain" or "connect a fact to his personal goals."

The difference between our categorization and Searle's reflects our

different interest. Austin and Searle are much interested the category of promises, for example, because promises have remarkable ethical effect; they put the speaker under an obligation. We class promises with other affective points, because generally the hearer will process promises in terms of his personal goals. Conversely, we treat "explanatory points" as a separate category because the information allows the hearer to reform his knowledge state in more reasonable terms. Austin and Searle consider these to be assertions, because they state truths about the world.

[Allen and Perrault 80] and [Perrault , Allen, and Cohen 78] have made a very sophisticated analysis of indirect requests in terms of planning and cooperation. However, they have so far restricted their analysis to what we have called need points and affective points; and, indeed, to conversations which open with the expression of a need point. It is not clear that their work can be extended to all the cases where it is important to analyze intention in utterances. In particular, a goal based analysis of empathetic or argument points, if possible, would involve rather subtle goals relating to the relationship between the speakers, or to the structure of the conversation, rather than the rather mundane goals (catch a train, open a door) considered by Allen, Perrault, and Cohen.

[Wilensky 80] distinguishes between external points, the goals of the story teller, and content points, interesting features of the plot structure. His research focusses on the latter, whereas we have examined

the former.

[Lehnert 78] discusses the question of determining the meaning of questions, and gives examples where scriptal knowledge and general inference rules associated with semantic primitives are needed to find the meaning of a question. We propose that this kind of information is not applied directly, but rather it is used to build up a model of the speaker; and that very similar processes are needed for all text understanding.

1.3 Overlaps in point categories

We have been talking about points and point categories thus far as though they are very clear-cut. Now let us look at a few examples which show that this is not always the case.

Sometimes it is difficult to say what THE point of a statement is. Consider the following exchange:

Person 1: I think Reagan is doing a good job on inflation.

Person 2: I don't. Did you hear the price of gasoline is supposed to go up to!$1.50 a gallon by next month?

One could say that Person 2's statement was an argument point, namely that inflation is getting worse. But assuming Person 1 drove a car and didn't know about the predicted gasoline price increase, the statement would be a good candidate for an affective point, also. The fact that either of these could be the point of the statement is evident in the way in which the conversation could proceed. It would not be surprising if Person 1 responded, "Really? I'm going to have to start taking the train to work from now on," reacting to Person 2's statement

as an affect point, or if he responded, "Yeah, but interest rates are going down," reacting to it as an argument point.

Either one of these two points would be reasonable for the understander to derive. There is no "correct" analysis. We are maintaining that deriving a point is a part of processing, specifically related to the choice of an "inference path". There is no "correct inference path". Understanders choose to process idiosyncratically. When an understander's choices do not correspond to the ones the speaker would have had him make, the speaker will feel that the understander has "missed the point". But that is a relative judgment, not an absolute one.

Sometimes, it is hard to say just how much the point contains. Consider the following example:

Person 1: The Yankees are going to do well this year. They just signed Dave Winfield.

Person 2: Good. The team needed a good leader.

Person 1: I was thinking more about his hitting power.

Here, Person 2 realizes that Person 1 mentioned the fact that the Yankees signed Winfield because it supports the assertion that they will do well this year. In this sense, he got the point. But Person 2 has found a different reason for the support than the one that Person 1 had in mind. Because of this, it still seems that he has missed at least part of Person 1's point.

So we see that points are more complicated than simply finding some assertion which falls into one of the several point categories. Which

category does the point really belong to? What does the point include? Sometimes it is difficult to answer these questions. We will not discuss these questions in detail. Rather, we will focus on the utility of the various categories. The "point" of points, and point categories, is to capture some generalizations about intentional communication. As we will see, these generalizations say things about what how the understander processes a point given its category, how to find a point of a given category, and also how to find the category of a given point. These generalizations are knowledge which we feel the understander must possess and use in order to process natural language.

1.4 Misunderstandings

Communication between speaker and hearer is not always perfect. Some cases of communication failure are trivial; for example if the hearer cannot hear what the speaker said. Other failures shed considerable light on our theory of points. It is worth noting that there is a special category of responses, such as "So what?" or "What's your point" which a hearer can use if he has parsed the input, but is unable to derive any point. In other cases, the hearer may mistake the point, and indicate this by an inappropriate response. The speaker will then presumably give him some more indications of how to process the statement. Some examples:

Example 1

A: The cat is scratching at the door.

B: Yes.

A: Why don't you let her in?

A intends his remark about the cat as a need point. B categorizes

the remark incorrectly as a point of information. This error might be made due to inattentiveness or simple rudeness; it might be made if B is ignorant of the ways of cats (he thinks it's simply sharpening its claws); it might be made if A's request is for some reason inappropriate (e.g. A is closer to the door). A makes his point clearer by rephrasing his remark as a question about the desired action.

Example 2

Holmes: Watson, do you remember seeing any cow-tracks today?

Watson: Yes, several . . .

Holmes: Well now, Watson, how many cows did you see on the moor?

Watson: I don't remember seeing any . . .

Holmes: Do you remember that the tracks were sometimes like that ... and sometimes like this ... and occasionally like this?

Watson: No, I cannot . . . What is your conclusion?

Holmes: Only that it is a remarkable cow which walks, canters, and gallops.

(The Adventure of the Priory School, Sir Arthur Conan Doyle)

Watson is unable to make the necessary connection between the facts about the cow tracks and any important feature of the case. From his knowledge of Holmes, he knows that Holmes' observations have an argument point, but he is unable to find the evidential connection. Hence he is baffled. His failure to make the necessary inferences is due to denseness, and perhaps also to ignorance of the tracks made by various gaits. Holmes clarifies himself by supplying the forward inference from the nature of the tracks to the nature of the gait, and leaves it to Watson to make the further inferences that the "cow" tracks were made by a horse.

Example 3

X: The young generation is incredibly degenerate. Did you know that Judy and Richard just had a child?

Y: Oh well, I've not sure that marriage is really all that important for happy family life.

X: But Judy is Richard's sister!

This is similar to the previous case, but somewhat more subtle (from a technical point of view.) Again, Y successfully categorizes X's statement as an argument point, and X (unlike Holmes) has made his overall point clear. Y is able to find a connection between X's evidence and his assertion, but owing to ignorance of the situation, he finds the wrong connection, thus missing the point. X clarifies the point by supplying him with the missing situation.

Example 4

Mrs. Brown: Suzy Jones is pregnant.

Mr. Brown: What do I care?

Mrs. Brown: Well, they use the same contraceptive method we do.

Note that it is entirely possible that Mr. Brown already knew about the Joneses' birth control method. He may merely have overlooked that source of relevance. This failure, either a failure of knowledge, or a failure to retrieve all the relevant information, led him to be unable to categorize the point of the utterance at all.

Example 5

P: I got a B in my computer course.

Q: Oh, congratulations!

P: No, no, I was expecting an A.

Q correctly realizes that he is being called upon to sympathize with P's emotional state, but he errs in judging what that state is. This is, again, due to a failure of knowledge. It is worth emphasizing

the subtlety of the interpersonal knowledge which may be needed to extract the point here.

Example 6

M: The Khmer Rouge have murdered two million Cambodians.

N: So what?

M: How can you be so callous?

This is an example of a speaker failure. M has made a mistake as to N's probable interests. His only approach is to try to awaken N's dead moral conscience.

There are sometimes ways of conveying the point which are essentially immune to misinterpretation. "Pass the salt" is one such. "The Khmer Rouge have murdered 2 million Cambodians" is another. If the hearer responds "So what?" there is nothing more to be said. However, people are rarely this direct; rather, they generally phrase what they say in ways that require their hearer to make some inferences. The most common reason is brevity. In example (3) above, Mrs. Brown could have said, "Suzy Jones uses the same kind of contraception that we do, and is pregnant. This shows that the method is fallible, and that maybe we should do something else." This would have eliminated any chance of misunderstanding, but it would have also taken twice as long to say. In (1), the indirectness is motivated by etiquette; in our society, it is more polite to word requests indirectly. In (2) Holmes is merely teasing Watson.

2. Points & the Speaker

We believe that our notion of points is a useful one in modeling both listeners and speakers in conversation. Although most of what we will have to say about the application of the concept of points to actual process modeling will concern the listener's end, there are some important things to be said about points and the speaker.

First, of what use are our notions of points and processing categories to someone wishing to build a conversational speaker? What we are proposing in effect is a stereotypical model of a listener in a conversational situation. That model views the listener as a processor with a small number of modes for dealing with input. Distinguishing these modes allows the speaker to store the following necessary information with each mode:

1. The typical results of processing an input in the mode, e.g., expectations about what happens when people are interested in or affected by something.

2. What sort of input will or is liable to be processed in the mode, e.g., how to decide what is interesting or what constitutes a usable explanation.

3. Knowledge of the appropriate way to express an input intended for processing in a given mode. For example, "Guess what!" is an appropriate preface to an interest point, but would sound strange preceding an explanatory or prescriptive point.

4. (more generally) Knowledge of what type of situation is appropriate for the expression of inputs intended for processing in a given mode. For example, explanatory points do not tend to make good cocktail party conversation, and overlooking this fact may cause the speaker to be perceived as being arrogant.

Thus our processing categories give the speaker a framework for storing necessary knowledge about the effects of communication on the

listener. This knowledge will be crucial for the speaker to construct concrete plans for fulfilling his conversational goals [Robertson and Johnson 81].

Consider an example where referencing this store of knowledge is used in a speaker's planning process. When one gives a talk or presentation, the goals involved are generally to provide some (hopefully) key explanation to the listeners, and at the same time impress them with one's prowess in the field. However, instrumental to these goals are the requirements that the listeners pay attention and be reasonably happy. A speaker faced with the necessity of insuring that these requirements are met will often make use of the fact that being interested in something (anything) tends to increase attentiveness and general happiness. This is an example of the type of knowledge referred to in (1), above.

To make use of this knowledge, the speaker will reference knowledge of type (2). He will attempt to discern what types of inputs will be interesting to his listeners. He may of course hope that his material is intrinsically interesting, but more specifically he may make use of the knowledge that people are generally more likely to be interested in what is familiar to them. Thus one way of attempting to interest an audience is to talk about their immediate environment - their college, or their hometown. This is a common practice of lecturers visiting a location. A more abstract version in academics is to try to relate one's research contributions to the particular interests of the

listeners (e.g. - "my new programming language would make writing natural language understanding programs much easier.")

Having found a concept he believes will be of interest to his audience, the speaker will now need knowledge of type (3) to be used in framing the exact way he wants to say it. This is particularly important in this case because the style of presentation can drastically affect the interest of the listener. Thus the speaker will try to use the knowledge he has stored here to heighten the interest of the audience as much as possible.

This example illustrates one particular use of our listener model; namely, the use of processing category knowledge to achieve an instrumental goal through communication. A similar example would be "buttering up" someone in order to achieve some other goal through them. The unifying feature in these cases is that they begin with a desired effect upon the listener, and proceed by reasoning backward from effect to processing category to input. We do not always follow this path in constructing our communications however.

Sometimes we begin with the goal simply of talking to someone. We may do this merely to be polite, or because we are bored, for example. In this case we will need knowledge of type (4). We will want to know, in other words, what type of thing that we could say would be appropriate to the situation. We may decide, for example, that it would be appropriate to discuss things which affect the listener. In this case we proceed to access knowledge of type (2), to determine what in

fact fulfills this requirement. (Note that the general knowledge of (2) will have to interact with specific knowledge of our listener. In fact, one form that knowledge of type (2) may take is principles for accessing one's knowledge of the specific listener in a useful way). From this point, of course, we apply knowledge from category (3) to produce our output.

We may make use of the knowledge of (4) in a converse way. Sometimes a situation makes it appropriate to communicate in a certain way with a listener. For example, if we know of a situation that will have a drastic effect on someone, and we know they are not aware of it, we will generally want to tell them about it. Likewise for some highly interesting fact. If we observe a child doing something which seems dangerous, such as playing with matches, we will often want to give them an explanatory or prescriptive input. These are all general rules which guide us to communicate with people in response to situations. Category (4) comprises these rules. Note that in applying them, we will have to make use of category (2) knowledge, to decide if, in fact, some fact will have the triggering effect on the listener.

Thus, we have suggested that a speaker will have rules leading to the generation of output which are based on the assumed processing categories of the listener, and that there are bodies of information specific to each category which are necessary to processes carried out in service of these rules, including the process of actually formulating a specific output. This amounts to saying that a speaker has a model of

the general listener which make reference to his processing categories and knowledge organized by these categories.

3. Points & Listener's Processing

We have looked at points in relationship to the speaker's task in communication; now let us consider the listener's conception of the point of the speaker. One important thing to note is that for the listener to get the point is not the same thing as for him to process the input in the manner the speaker intended. That is, the speaker may not succeed in producing the desired effect upon the listener, but this fact does not imply that the listener does not realize what that desired effect was.

We have already looked at an example where this fact was implicit. We noted that a reasonable response to "the Yankees just signed Dave Winfield" might be "I don't like baseball." To generate such a response, the listener would have to be aware that the speaker intended the statement to be processed in terms of interest, despite the fact that he did not find it interesting. This implies that the listener must be processing the input not only by determining what processing category is appropriate for dealing with it, but also in terms of what processing category the speaker intended the listener to choose.

In the Winfield example the understander failed to derive a processing category for the input, but was nonetheless able to deduce the intended processing category. Alternatively, it is possible find a way to process the input without knowing what the intended processing category was. For example, if someone inadvertently gives you some information concerning something which, unknown to them, is very

important to you, you are very likely to process the input in terms of its effect on you. However, the fact that you are affected by this input does not imply that the speaker's point was necessarily to affect you. The intended processing category may be unrelated, or unclear. The speaker may have had no idea that you might be affected by the particular information involved. Furthermore, this fact may be entirely clear to you, since it is quite possible that you could know that the speaker wouldn't have the knowledge to infer that you would be affected.

A simple example of how this could happen might be in a situation where the speaker is using facts to support an argument point. Say he is making a point about the shocking number of children born out of wedlock these days, and he mentions as an example that Mary Jones is pregnant. If you have been secretly dating Mary, this fact may well be of tremendous import to you, and you are likely to process the statement in exactly the same way you would had it been staightforwardly intended as an affect point by someone who was aware of your relationship with Mary. However, this in no way implies that the speaker's point had anything to do with that. Furthermore, the fact that this was not the speaker's point should be perfectly obvious to you.

These examples make it clear that two separate but related processes are occurring in the listener when an input is being processed. The listener is both deciding on how to process the input and trying to determine how the speaker intended him to do so. In order to make use of the notion of points in a process model of an

understander, we will have to make clear how these two processes interrelate.

The relationship between the process of tracking points and the process of deciding what to do with input is not a simple one, since either process can potentially be helpful for the other one. Deciding that an input is interesting, or inferring from it that the person you are talking to needs help, is not equivalent to determining that that was the person's point, but it is certainly good evidence to that effect. Likewise, knowing that a person has an explanatory point does not ensure that you will grasp it, but it can at least help to guide you in the right direction for processing.

In general, whenever we decide that an input belongs in a certain processing category, we check to see if that might have been the speaker's point, and whenever we decide on the speaker's point, we see if we can process the input in the manner he intended. These two situations may be seen as roughly constituting a distinction between "bottom-up" and "top-down" derivation of the point. That is, if we use some cues taken from our knowledge of the speaker or the situation to decide what sort of point the speaker probably has, and use that information to direct future processing, we have analyzed the input top-down by making heavy use of our expectations. On the other hand, if we can't figure out from the situation what type of point of the person has, we may have to first decide that the input fits one of our processing categories by attempting to apply each of them (is this

interesting? does this affect me? etc.). Having found a processing category which is appropriate, we may then ask ourselves "was that his point?" In this case we have done our processing bottom-up, being guided by our analysis of the input rather than our expectations.

Let us look at an example which makes this distinction clear. Consider the following case: an official says to you, "You are standing on the grass." Accordingly, you move off the grass. Now, you may have known perfectly well that you were standing on the grass, and you may even have known that the official knew it. Therefore, your action was not prompted by any new information, as such. Why, then, did you decide to move?

The answer is that the official's point was not an informative point, it was a need point (in a broad sense). The official was requesting that you move off the grass. There are at least two ways in which you could determine the category of his point. You might notice that you already knew the content of his statement, and therefore the statement presumably served some other function. It is a general rule that need points are often presented as statements of the problem (compare "I'm hungry", "It's cold in here", etc.) Therefore, your being on the grass constitutes a problem for the official, which can be remedied by appropriate action on your part.

Alternatively, you may use your knowledge about officials. Usually, when an official addresses you, it is with the purpose of giving you an order. Frequently, the owners of grass prefer that you

stay off it, and they employ an official to enforce this preference. The understander may, by a judicious application of these rules, arrive at an interpretation of the officials remarks.

Actually, both methods must be available to the understander. The former must be used if someone who does not look like an official tells you, "You are standing on the grass." The latter must be used if an official shouts "Hey you!" at you, and you wish to respond before he gets angrier.

There are, in fact, a number of studies in the psychological literature of the processing of indirect requests, such as the one above. The results of Gibbs (1979) suggest that, when it is available, contextual information is heavily used in determining that a statement is a request. Specifically, he showed that people normally do not need to first assess the "literal meaning" of an utterance, and then realize that the indirect meaning is the intended meaning. Instead, subjects seem to be able to correctly uncover the intended indirect request, without having to first understand a "literal meaning". Gibbs found that in a story context, indirect requests are understood faster than the same sentences used as literal questions.

The meaning of such results can be made clear in the context of our theory. When we use our knowledge of officials to understand the official's point in the above example, this is a case of almost pure top-down processing on our part. We do not have to derive the literal meaning of the sentence because the situation, and perhaps some

meta-linguistic cues such as the official's tone of voice, allow us to realize the type of point he has without reference to the content of the input.

On the other hand, the alternative method of processing is almost purely bottom-up. Here we derive the content of the input, recognize that the input seems to provide information, and ask ourselves, "could that have been his point?" Because of the obviousness of the information, we are forced to answer "no" to this question, and it is the subsequent search for what the person's point might have been which produces the realization that the point is a need point.

We will now consider how to use the two processes of finding the speaker's point and finding a processing category to build a model of a conversational listener. To begin with, we can recognize the following principles by which the model should operate:

1. For each input, at least one speaker's point must be derived. If this is not the case, understanding has failed.

2. For each speaker's point derived, the listener must consider whether the input can be processed in the intended way.

3. For each input, any number of processing categories may be applied (depending on the importance to the understander).

4. For each processing category applied, the understander must consider whether the speaker may have intended the application of that category.

5. The input content must be understood at least sufficiently to decide what the speaker's point is and to determine if the listener can process the input in the intended way. If one of these cannot be accomplished, and the input has not been sufficiently understood to make it reasonable certain that it will not be accomplished, the process must continue. The process of understanding the content must also continue if a processing) category is chosen, processing in the direction indicated by that choice.

Based on these principles we can construct the following algorithm:

Is there enough a priori and situational knowledge available to guess the category of the speaker's point?

If yes,
- assume that the speaker's point is in that category
- see if the input can be processed in that way
- if so, do and STOP
- otherwise, continue

Try to find a processing category applicable to the input [i.e. - ask "does this interest me?" "does this affect me?" etc.]

Was a category found?

If no, ask "what's your point?" and STOP

Process the input according to the category

Could this have been the speaker's point?

If yes, assume this is the speaker's point and STOP

Return to beginning [now considering the fact that we derived a category which couldn't have been the speaker's point as part of the situational information]

We have oversimplified somewhat here in that we have neglected to say anything about how this processes relates to other processes involved in understanding. In particular, we may think of the type of understanding processes upon which previous research in natural language processing has typically focused: representing text and forming causal and knowledge structure based (eg - scripts, MOPs) relationships between the input and other knowledge. The relationship between understanding and point derivation is a more complex sort of process than the one we have been looking at.

For one thing, we can understand an input, at least to some extent, without having the slightest grasp of the speaker's point in saying it,

and without any applicable processing category for it. That is, someone could come up to you and make almost any sort of simple, declarative statement - say, "John and Mary just got married" - and you can process the content of the input, in this case by instantiating the knowledge structure associated with "wedding". However, this input (assuming you don't know John and Mary) will not fall into any of your processing categories, and the point of the speaker will be completely opaque to you. In fact, just because you could process such a sentence, it doesn't necessarily follow that you would, given a lack of understanding of the point of the processing.

On the other hand, there are times when the search for the point of a statement results in true understanding. This is the case in the "standing on the grass" example above. Only by considering the intent of the speaker in affecting you can you actually derive the meaning of the official's statement.

Furthermore, as was the case with that example, one can at times anticipate the point before the content of the communication is actually present. However, there are also times when the point is grasped only after lengthy consideration of the content.

So we can see that the search for the point may fail, though understanding proceeds; that understanding may, on the contrary, depend on the search for the point; that the search for the point may proceed the processing of the content, but may also follow it.

The point here is that it just does not seem possible to specify

any sort of serial relationship between the search for a point and general understanding of the input. It seems most reasonable to think of the point processor as a routine that acts in some sense as a watchdog over the understanding process, serving to monitor whether that process is producing reasonable representations of the input, to guide the attention of the understanding process away from input which seems to lack a point, and to direct further processing of an input for which a point has been found.

Specifying this relationship algorithmically is certainly a complex task, and it lies beyond our scope here. However, we can give some principles for the interaction of the process of deriving the point and the rest of understanding. In particular, we can assert the following:

1. If there is some anticipation of the content of the input, this will not occur before the application of some top-down knowledge of the speaker's likely point. That is, it is hard to imagine processes which give us a reason to guess at the content of the input without doing so by way of a consideration of the speaker's intent. This restriction is not inviolable, but at the least it is not likely to be profitable for the understander to try to guess the content before considering the intent of the speaker, and thus the likely point category for the input.

2. The process of determining the applicability of processing categories cannot begin until at least some understanding of the content has occurred. Whether the process is triggered by the derivation of a possible point, or simply begins of its own accord, it has nothing to operate on without some understanding of the content of the input. However, this understanding must by no means be complete, as the example above where application of a processing category begins before the input is even received demonstrates.

3. The effect of choosing a processing category will be to direct further processing of the input content (with respect to the structures in memory). Therefore, once the process for choosing a processing category has chosen, the processing of the content must continue along the direction specified by that choice.

4. Tracking and Deriving Points

4.1 Tracking the Point

In the previous section we argued that keeping track of the speaker's points is an important part of an understander's processing of input. In this section we will discuss in more detail just how this process is helpful and what the processing involved looks like.

Why is the consideration of the speaker's point advantageous to the listener? There are three basic reasons:

1. Understanding the point helps in understanding and processing the input by allowing the listener to reason teleologically about the input
2. Understanding the point provides a minimal criterion for the decision on when sufficient processing of an input has occurred
3. Understanding the point helps the listener decide how to respond appropriately

Let us consider each of these in more detail.

Teleology

To reason about something teleologically is to consider it in terms of its ultimate purpose. Often knowing the purpose for which something is intended is an extremely important clue in discovering other things about it, as knowing the purpose of a piece of machinery might help us to understand its mechanism. This type of reasoning from purpose can be similarly useful in understanding any sort of intentional behavior.

In communication, the purpose of a speaker is captured in the point of what he is saying -- at least, at the superficial level (such phenomena as "lying" consist of a purpose at a deeper level which might

also be of use in understanding the statement). It follows, then, that an ability to grasp the point of a speaker is necessary to enable this type of reasoning about the speaker's utterances.

How might reasoning of this sort aid us specifically? Let us look at an example from [Schank 81]. In 1979, when President Carter was discussing the Soviet Union's invasion of Afghanistan at a news conference, he alluded to the Munich conference of 1938. When the President of the United States addresses the people, he in principle can have only one point, by virtue of the nature of his role. His point has to be: "the situation I am describing can affect you (probably adversely) and something must be done about it."

This knowledge allows us to infer that point will involve the *negative impact the Soviet Union's activities* in Afghanistan will have on us. Knowing this to be the case, we immediately know some things about his comparison with the Munich conference. We know that the goal of the comparison will be to make his point, so the comparison must concern the negative aspects of the Afghanistan situation and the thrust of the example of Munich will probably be to draw out the nature of the possible negative consequences to us.

Suppose that we don't know exactly what the Munich conference in 1938 was about. Our reasoning above gives us some leverage in guessing what the subject was. If we know that Germany invaded Czechoslovakia at around that time, we can hypothesize that this was the subject of the conference, and that Carter means to point out the negative consequences

which followed from that action, those being, of course, World War Two. On the other hand, were a "1938 conference at Munich" mentioned in any other context, we might have little or no idea what it was about. Here we use our knowledge of the speakers's point to make a reasonable attempt at understanding something we might not otherwise have gotten.

Criterion

As we said in the previous section, tracking the speaker's point provides a kind of check on the process of communication and our understanding. In normal situations we minimally expect that we will be able to figure out what a speaker meant for us to do with his input, whether or not we can actually make any use of it. If we cannot even accomplish this, then we assume that the communication is faulty. It is in this situation that we normally ask "what's your point".

The key to why tracking points is particularly important in monitoring the success of the communication is that it is in general a weaker constraint that we should be able to guess a point than that we should be able to apply a processing category to the input. It is easier to figure out that a person meant an input to interest or affect us than that it is to actually be interested or affected, in the same sense that it is more likely that we can figure out that someone is joking than it is that we will be amused. It is possible to know that someone is arguing with you or expressing need to you without being able to make much sense of what they are saying. The reason for this is that it is much easier to apply top-down knowledge to the guessing of a person's point than it is to decide what processing category an input

belongs in. In fact, most of the top-down knowledge which can be applied to the latter task comes by way of the former - for example, top-down knowledge which might be of aid to you in processing what a teacher tells you in class is most likely to come from your expectations about what he will want to tell you, in other words, from knowledge concerning what his point is likely to be.

Response

In conversation, each type of point we have mentioned will have appropriate and inappropriate responses. The appropriateness of a given response will be governed by a number of factors, but one of the most important of these is the point the speaker intended.

For example, suppose someone tells you, "Reagan was just elected president". An appropriate response might be "This country is really swinging towards the right." However, "Gee, that means he'll be moving out of his house" seems rather inappropriate. Nonetheless, given that Reagan was your next-door neighbor in California this might be a reasonable reaction to the statement. For the response to seem apropriate, though, the speaker would have had to have as his intention making some point about the effect of Reagan's election on you via this connection. Thus the appropriateness of the response in this situation depends not on how reasonable a given reaction to the input is, but rather on the relevance of the response to the speaker's point. The only way for Reagan's next-door neighbor to avoid appearing entirely egocentric in conversation is for him to be able to reason about what the speaker's point is.

4.2 Deriving the Point

Each type of point is derived and tracked by its own set of procedures. In this section, we hope to highlight some of these procedures.

Affective Points

Since affective points deal with plans and goals, the first step which an understander must take in deriving an affective point from a statement is determining which of his plans and goals might be relevant to the statement. The goals and plans which we will refer to here are discussed in detail in [Schank and Abelson 1977].

Often certain individuals, places, or objects have certain plans or goals associated with them. For example, one's spouse is associated with the goal S-SEX, the bank is associated with D-CONT(money), and a car is associated with D-PROX. Therefore, one way of finding relevant goals is to find those goals associated with the people, places, and objects which are mentioned in the input statement.

Events also have specific plans or goals associated with them. For example, hitting someone might be associated with the planbox:OVERPOWER, and eating is associated with the goal S-HUNGER. So the events mentioned in a statement are also a good source of relevant goals and plans.

Sometimes, the actual goal or plan which the understander has may not be referred to in the point of a statement, but instead a related goal or plan may be mentioned. For instance, if Person A says, "I want

to rob a bank," and Person B replies, "I have a gun I can lend you," Person B is addressing the OVERPOWER planbox, which is part of a plan which would satisfy Person A's goal of robbing the bank. Therefore, another way for the understander to determine which of his plans and goals are relevant to a statement is to try to find a relationship between a goal or plan mentioned in the statement and his own goals and plans.

Once relevant goals are found, the process of determining if a goal is affected by an input statement is a fairly simple one. If the statement itself does not affect the goal, then we predict some possible outcomes from the statement and see if they affect the goal. But we need not go any further than this in our examination. We do not look for convoluted paths that will show how something we have just heard affects us. It is worthwhile to point out however that some people do precisely such convoluted derivations. We usually refer to such people as paranoid. The difference between paranoid and normal may simply be, with respect to the issue of points anyway, just a difference of degree of tracking of affective consequences.

Empathetic Points

Often in conversations, one may begin to question why a speaker is talking about a particular situation in which he is involved. It is assumed that the speaker is not merely making an interest point, since discussing personal experiences is not typically done solely for the interest value. On the other hand, the listener may not be able to figure out a way in which he could possibly be of help in the situation,

so it does not seem that the speaker is making a need point. People often seem to discuss situations that they may not be able to do anything about (at least for the moment). Further, the situation does not seem likely to affect the speaker directly, so it is not a typical affective point. Rather, it seems that the person is talking about something that affects them simply so as to make someone else understand how they are feeling.

In cases such as these, where a speaker is relating something of personal import, and the possibility of need and affective points have been discounted, the listener assumes that the speaker is making an empathetic point. The speaker wants to feel that we truly understand how he is feeling. To do this, one tries to imagine an analagous situation affecting one's own life. Thus, the listener processes the input as if he were the focus in the situation the speaker is describing. Inferences then proceed along the path of determining how he would feel, which he then try to relate back to the speaker's situation. This enables him to make the appropriate inferences necessary to understand how the speaker is feeling. Unless one has grasped that the speaker is making an empathetic point, he might not be inclined to make those inferences to that extent.

Need Points

Deriving need points s largely a function of pursuing two questions:

1. Is the speaker expressing a difficulty or problem?
2. If so, is it likely that you can assist in providing a resolution to the problem?

In cases in which the speaker is not asking for help directly, the first question is pursued by tracking his goals and plans. People often describe their problems by discussing their past planning failures. "I've tried to work out all of the problem sets, but I still flunk the tests." In other cases, people express their problems by stating plan precondition failures. "I want to ask out Farrah, but I don't know her phone number." The general problem of following a person's goals and plans is very complicated. [Wilensky 78] Note that the same problem comes up in deriving empathetic points.

The second question is the key to recognizing the difference between a need point and an empathetic point. A speaker is probably asking for help in cases in which the listener is just the person that can provide it. One way of recognizing this is via role stereotypes. Consider the example of the failing algebra student. The listening student is likely to interpret the academic lament as a need point only if he has a reputation for doing well in algebra. In general, people that have been stereotyped as good students tend to interpret their fellow students' academic complaints as expressions of need points. On the other hand, if the listening student has a reputation as a poor student, he is unlikely to interpret the other student's complaint as any more than an empathetic point.

Of course, the situation is not always this complicated. When problems are described as plan precondition failures, a listener need only determine if it is possible for him to do anything about the

precondition state. If you are Farrah's friend or relative, then you can easily provide her phone number. Of course, detecting a need point and deciding how to further process it are not always the same problems.

Suppose for example that the good student is heavily involved in a sport and does not have time to tutor his friend. Rather than taking the time to fully understand his friend's problem, he may elect to ignore the complaint as a request for help completely. Instead, he may decide to process it as an expression of an empathetic point. In this case, he might respond with something like: "Gee, that's tough." This demonstrates that someone may detect the real point of a speaker's statement, but choose to process it as if it were another kind of point all together. This tactic is a very basic kind of ploy in conversation that allows us to avoid direct confrontation. It is often considered impolite to phrase one's argument as a direct request because it limits the understander's possible replies.

Explanatory Points

Generally when an explanatory point is made, there is a great deal of contextual information which both indicates that an explanatory point will be made, and also indicates the matter to be explained. Frequently, of course, explanations are given in answer to direct questions. Explanations are given in lectures and in expository texts. Here, the context dictates that what is said will be an explanation; and the lecturer or author is obliged to mention the explanand before giving the explanation. (Overly formal mathematical texts, which present results before motivating them, are counter-examples here; but these are

notoriously difficult to understand.) If a man is obliged to explain his actions, he can rely on the fact that his audience knows what he has done and are expecting an explanation. There are exceptional cases where an explanation is offered without the context demanding one, e.g.

A : Dammit, the banks are closed today.

B : Yes, it's Lincoln's birthday.

However, it is difficult to think of cases where explanations are offered where the matter to be explained has not been brought up.

The processing of an explanation and the question of what constitutes a valid explanation are difficult problems beyond the scope of this paper.

Prescriptive Points

One obvious cue for deriving prescriptive points top-down is to consider the speaker. Some people are far more likely to prescribe things than others. Likely people are parents, clergymen, and doctors (of the type who don't believe in explanations). In fact, anyone who considers themself to be an expert at something which they don't consider you to be an expert at is likely to start tossing prescriptive points around when that thing becomes the topic. Thus people collect such advice as "buy low, sell short" or "always downshift when braking your car" with little or no notion of what these things will do for them.

Prescriptive points are often signaled by linguistic cues: "let me give you some advice" is a common one. Any imperative which is not directing a specific action (such as a "never" or "always"

construction") almost certainly has a prescriptive point. Situational information (other than the topic of conversation, mentioned above) may come into play occasionally. For example, having some sort of problem, or mentioning one that you are having, will often provoke a prescriptive point directed at preventing you from having any more such problems in the future.

Argument Points

Obviously the most common source of top-down knowledge that an argument point is coming is the knowledge that an argument is going on. A slightly less trivial case is when we can guess that an argument is about to start. In general, we can expect that if we contradict a position we know someone to believe strongly, or contradict something they just said, an argument point will be forethcoming. We can additionally expect this to be the case when a person makes some sort of general claim, like "Reagan is an idiot!" In general people do not make such statements without following them with evidence to support their position. A statement such as the one above, that simply stops there, would seem pointless.

Thus, top down derivation of argument points will generally stem from an understanding of the situation at hand, although knowledge of the person involved will also be useful. In addition some topics, such as politics or religion, have reputations for starting arguments, and may thus be factored into top-down consideration of points.

In deriving an argument point bottom-up, it is necessary to do forward inferencing. If a statement can trigger some type of inference

that supports a proposition which has either been in the conversation already, or is known to be important to the speaker, it may be interpreted as an argument point. The search for such inferences may be guided by the topic of the conversation, since if an argument point occurs in the middle of the conversation it will generally be as evidence for a proposition which is either the topic of the conversation or closely related to it.

Deriving Interest Points

The category of interest points is a weak one, in that categorizing a statement as an interest point tells one very little about how to process it. Interest points are merely to be added to one's stock of knowledge about the subject in question. For this reason, context is generally unimportant in determining the point of an interesting statement. It matters very little whether "The Yankees are trying to sign Dave Winfield" is spoken by a debutante or a policeman, or is written in a newspaper; the point you derive is the same. The weakness of the category suggests that it is the default category, to be used also if all other categories fail. For example, if you have had an argument with someone on some subject, it will be very difficult in subsequent conversations to introduce any information on that subject as an interest point. Your conversant will almost always assume that you are trying to make an argument point, since this is the stronger category.

Despite this, the category of interest points is an important one, if only because it represents a very common form of understanding. A

statement may be understood by connecting it to one's interests; and once one has made the connection, one need not worry whether the speaker meant a need point, an affect point, etc. Moreover, the categorization will lead you to a particular set of questions - "Why should that interest me?", "How could he imagine that that would interest me?" etc. - which will be helpful in formulating an appropriate response.

Categorizing a point as an interest point is also useful when the point is obliquely worded. If someone says to you "Dave Winfield flew into New York last night", you may realize that he has some interest point about baseball to make, and this may guide you to the deduction that Winfield is talking to the Yankees. This explains the anomalous nature of the following:

A: Dave Winfield flew into New York last night.

B: The Yankees are trying to get him?

A: No, it was his mother's birthday.

5. A Detailed Example

We will now examine an example in some detail which will show that deriving points is necessary for proper understanding. Consider the following conversation between a husband and wife:

Wife: I went to the doctor's office today.

Husband: Are you sick?

Wife: No, I'm fine.

Husband: Then why did you go to the doctor's?

Wife: Do you know that extra room we were thinking of adding on?

Husband: What about it?

Wife: Well, the doctor says we're going to need it.

The point of the wife's sequence of statements is clearly, "I'm pregnant." Let us examine how the husband derives this point.

The wife's first statement strongly suggests to the husband a category for her point, as well as fairly specific limitations on its content. A visit to the doctor usually implies illness, and this means that the wife's goal of preserving her health is being threatened. This implies an affective point, since affective points are transitive across relationships, as we said earlier. Therefore, the husband guesses that his wife's point is that she is sick. Since there is some doubt, however, about what exactly her point is, his reply of "Are you sick?" is a way of asking for verification that he has derived the correct point.

There is a problem here. Although the wife's first statement does suggest the point that she is sick, it does not explicitly state it.

There are other possibilities (as the husband eventually discovers). Why, then, does the husband make the inference that her point is that she is ill?

The reason is that the connection between going to the doctors' and illness is very strong. Therefore, the possible affective point here forces the husband to leap to a conclusion in attempting to find the point. In a less clear or less affective situation, the husband might very well have reacted differently to his wife's statement. For instance, drug stores are also associated with illness, but if the wife had said, "I went to the drug store today," this would not have created as strong an expectation for an affect point. In this case, the husband would have been more likely to respond with, "So what?" or something else meaning, "What's your point?"

Continuing with the present example, after the husband asks for verification that he has derived the correct point, she responds by saying, "No, I'm fine" or, in other words, "No, you don't have the correct point." At this point, he does not know what his wife's point is. Now he is justified in saying something which means, "What's your point?". (It is important to realize that his statement, "Then why did you go to the doctor's?", is an encoding of "What's your point?". To express our processing goals, such as finding the point, we often encode them using the context of the conversation. Just because we mean "What's your point?" we don't constantly say it.)

At this point, although the husband does not know the category or

the content of his wife's point, he does still have some expectations about what the point will be. He knows that her point eventually has to relate back to a trip to the doctor's. In other words, he is still assuming that she did have a point to her first statement.

There is a problem here which is analagous to the problem discussed earlier. Why does the husband still assume that his wife's first statement had a point? In many similar cases, it would be correct to conclude that there is no point to be found. Let us illustrate this with another example. In this example, the husband has been cheating on his wife with his wife's best friend, Ellen. Ellen is mad at the husband, and the husband fears that Ellen will tell his wife about them.

Wife: I got a call from Ellen today.

Husband: What did you talk about?

Wife: Nothing very interesting.

Husband: What else did you do today?

This example is very similar to the first example, in that the wife's first line sets up an expectation for an affective point, namely that Ellen told her about their affair. He asks a question, trying to confirm this expectation. Again, the wife does not confirm these expectations ("Nothing very interesting"), and apparently the husband's expectations for an affective point have been violated. As in the first example, the husband is now left with almost no expectations about what her point might be. But this time, he does not ask what her point is. Instead, he changes the subject, glad that she does not have a point.

So how can we explain why the husband assumes that there is still a

point in the doctor example? The reason is that the husband cannot assign a motive to his wife's trip to the doctor. In the affair example, he can understand why Ellen had called his wife: Ellen and his wife are friends, and friends call each other on the phone just to talk. But people do not go to doctor's offices just to talk. Because of this, the husband is not satisfied with finding an interest point; he wants an affective point.

Let us continue with the example. We have established that the husband is still expecting a point about the doctor's office. When the wife says that the doctor thinks that they will need an extra room, the husband tries to find a point which involves a doctor, does not involve illness, and which has something to do with an additional room. He is supposed to conclude that *his wife is going to have a baby*. Exactly how he does this is not really relevant to points. What is relevant, though, is that it is his searching for a point which leads him to look for a connection between these things, thus enabling him to infer that his wife is pregnant.

Searching for points is vital to inference-making. To demonstrate this further, let us suppose that the doctor is a good friend of the wife's. In this case, when the husband asked, "What's wrong with you?" and his wife told him that nothing was wrong, he could have inferred another motive to her visit to the doctor's: she was paying him a social visit. Then, if she had gone on to say that the doctor thought they needed an addition, he might have missed the point, since he was no

longer looking for one. He could have come up with another explanation as to why the doctor would say that they needed an addition, such as maybe their house would be worth a lot more if they built an addition, so it would be a good investment, and the doctor was giving them advice as a friend.

Notice also that the husband's search for the point not only explains how he was able to understand such an obscure reference to pregnancy, it also explains why he said what he did during the course of the conversation. His expectations for finding a point guided his questioning. So we see that this example illustrates two uses for points -- explaining how obscure inferences can be made in some situations, and explaining, at least in part, why a conversation proceeds in the manner in which it does.

6. Conclusion

Our examination of points has led us to the following conclusions:

1) Typically, a speaker will have in mind the way in which he wishes his listener to process an utterance.

2) Besides processing an utterance for content, a listener will also try to determine how the speaker intended him to process it. This information may aid him in understanding the utterance and in framing an appropriate reply. The processes used to determine the speaker's point involve a mixture of top-down and bottom-up reasoning.

3) Deriving the point is an important criterion for understanding. Therefore, processing must continue until the understander has satisfied himself, either that he has found the point, or that finding the point is not worth the effort.

4) Varieties of points may be profitably categorized, as above, by the type of processing which the speaker desires in the listener.

5) Restricting inferences to those which help find the point protects the inference process from combinatorial explosion. Points guide inferencing and also guide conversation.

What does this mean for AI programs? In the past, natural language programs have generally interpreted input and produced output as if in a vacuum, with no model (or at best a fixed, invariant model) of why the input came, and what the reason had to say about subsequent processing. Knowing what a speaker may have in mind, or knowing that you don't know

what he is getting at, is critical to understanding.

We are talking about making programs purposive. Just finding new inferences isn't good enough. A program must know if what it has found is satisfactory. Our definition of *satisfaction* must employ a concept of a point such that we could know when we have found it and also know to keep on trying.

REFERENCES

[Abelson 80] Abelson, Robert.
The Psychological Status of the Script Concept.
Technical Report 2, Yale University. Cognitive Science Program, 1980.

[Allen and Perrault 80]
Allen, J.F. and Perrault, J.R.
Analyzing intention in utterances.
Artificial Intelligence 15(3):143-178, 1980.

[Austin 62] Austin, J. L.
How To Do Things With Words.
Oxford University Press, New York, 1962.

[Carbonell 79] Carbonell , Jaime G.
Subjective Understanding: Computer Models of Belief Systems.
Technical Report 150, Yale University. Department of Computer Science, 1979.

[Cullingford 78]
Cullingford, R. E.
Script Application: Computer Understanding of Newspaper Stories.
Technical Report 116, Yale University. Department of Computer Science, 1978.

[Flowers, McGuire, and Birnbaum 81]
Flowers, M., McGuire, R., and Birnbaum, L.
Adversary Arguments and the Logic of Personal Attacks.
In Lehnert and Ringle, Computers and Natural Language, . Lawrence Erlbaum Associates, 1981.

[Grice 75] Grice, H.P.
Logic and Conversation.
In Davidson and Harman, The Logic of Grammar, . Dickenson Publishing Company, 1975.

[Lehnert 78] Lehnert, Wendy G.
The Process of Question Answering.
Lawrence Erlbaum, Hillsdale, New Jersy, 1978.

[Morgan 78] Morgan, Jerry L.
Toward a Rational Model of Discourse Comprehension.
In TINLAP-2: Theoretical Issues in Latural Language Processing-2. ACM, University of Illinois at Urbana-Champaign, July, 1978.

[Perrault , Allen, and Cohen 78]
Perrault, C.R. Allen, J.F. and Cohen P.R.
Speech acts as a basis for understanding dialogue coherence.
In TINLAP-2: Theoretical Issues in Latural Language Processing-2. ACM, University of Illinois at Urbana-Champaign, July, 1978.

[Robertson and Johnson 81]
Robertson, Scott P. and Johnson, Pete N.
The Interaction of Conversational Goals in Producing Statements in a Natural Dialog.
In Paper submitted to the Seventh International Joint Conference On Artificial Intelligence. IJCAI, Vancouver, B.C., Canada, August, 1981.

[Schank and Abelson 77]
Schank, Roger and Abelson, Robert.
Scripts, Plans, Goals, and Understanding.
Lawrence Erlbaum Associates, Hillsdale, New Jersey, 1977.
The Artificial Intelligence Series.

[Schank 78] Schank, R.C.
Interestingness: Controlling Inferences.
Technical Report 145, Yale University. Department of Computer Science, 1978.

[Schank 81] Schank, Roger.
Dynamic Memory: A Theory of Learning in Computers and People.
Lawrence Erlbaum Associates, Hillsdale, New Jersey, 1981.
in press.

[Searle 69] Searle, J. R.
Speech Acts.
Cambridge University Press, Cambridge, 1969.

[Wilensky 78] Wilensky, Robert.
Understanding Goal-Based Stories.
Technical Report 140, Yale University. Department of Computer Science , 1978.

[Wilensky 80] Wilensky, Robert.
Points: A Theory of Story Content.
Technical Report M80/17, University of California, Berkeley. Electronics Research Laboratory, 1980.

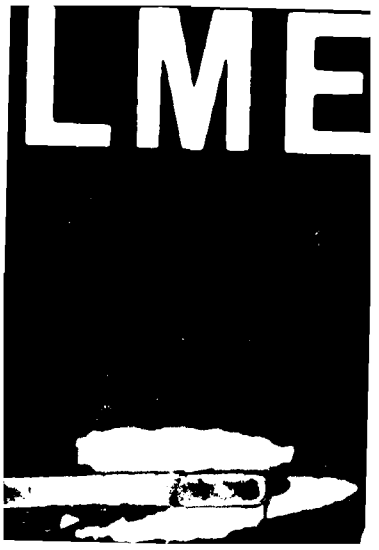
LME